



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Simulation of Drift-Compression for Heavy-Ion-Fusion

W. M. Sharp, J. J. Barnard, D. P. Grote, C. M.
Celata, S. S. Yu

March 18, 2005

Heavy Ion Fusion Symposium
Princeton, NJ, United States
June 7, 2004 through June 11, 2004

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Simulation of Drift-Compression for Heavy-Ion-Fusion

W. M. Sharp, J. J. Barnard, and D. P. Grote,
Lawrence Livermore National Laboratory, Livermore, CA 94550, USA
C. M. Celata and S. S. Yu
Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

Abstract

Lengthwise compression of space-charge-dominated beams is needed to obtain the high input power required for heavy-ion fusion. The “drift-compression” scenario studied here first applies a head-to-tail velocity variation with the beam tail moving faster than the head. As the beam drifts, the longitudinal space-charge field slows compression, leaving the beam nearly monoenergetic as it enters the final-focus magnets. This paper presents initial work to model this compression scenario. Fluid and particle simulations are compared, and several strategies for setting up the compression schedule are discussed.

I. INTRODUCTION

Beams for heavy-ion fusion (HIF) must be compressed lengthwise by a factor of more than ten between an induction accelerator and the final-focus magnets. The compression scenario favored by the US HIF program is to impose a head-to-tail velocity increase or “tilt”, so the beam tail approaches the head in a “drift-compression” section. The beam current and velocity must be accurately tailored before drift-compression in order that the longitudinal space-charge field removes the velocity tilt just as the beam traverses the final-focus lattice. Transverse focusing in the drift-compression lattice must also be carefully designed to ensure that all parts of the beam remain approximately matched as the beam expands to the larger radius needed for final focusing.

An important problem posed by drift-compression is how to prepare the beam velocity and current profiles before compression. Early work by Ho, Brandon, and Lee [1] used analysis and one-dimensional numerical simulations to model a beam compression sequence that produced uniform current and velocity on target. The effects of longitudinal space charge were ignored during shaping, on the grounds that shaping could be done in a sufficiently short time. A later paper by Sharp *et al.* [2] developed a more elaborate scheme where longitudinal shaping fields transform the input beam profile into the form found by allowing a compressed beam to expand backward through the drift-compression lattice. This shaping method was tested using a three-dimensional (3-D) fluid/envelope beam simulation that had a simple self-consistent model of the longitudinal space-charge field. Also, recent work by Qin *et al.* [3] uses an elegant three-dimensional envelope approach to develop several scenarios for self-similar compression.

This paper addresses two questions: whether the fluid models in earlier work adequately capture the

physics of drift-compression and what algorithm might be used to shape the beam current and velocity profiles of HIF beams before longitudinal compression. Other critical questions, such as how much the total emittance grows, whether a beam halo develops, and how these processes scale with beam and lattice parameters, will be addressed elsewhere. Section II describes three models of beam compression: a cold-fluid analytic description, a fluid-like simulation, and a particle simulation. In Section III.A, we compare results of these models for an idealized compression scenario, and Section III.B uses the simulations to demonstrate several methods to prepare a beam for drift compression. We conclude with a brief summary of our findings.

II. COMPRESSION MODELS

Three models are used here to describe beam compression. We compare the results of a simple cold-fluid analytic model with numerical cases from a fluid/envelope code CIRCE [4] and from a three-dimensional (3-D) particle-simulation code WARP3d [5]. Two important measures of beam compression are the ratio of the initial length or duration to the value at stagnation, called the compression ratio, and the distance the beam midpoint propagates before stagnation, termed either the compression length or stagnation distance. In addition, the numerical models provide information about beam dynamics, such as the radial envelope and, for the WARP cases, transverse and longitudinal emittance.

With several simplifying assumptions, a simple cold-fluid model provides usable estimates of the compression ratio and length. The critical simplification is to model the longitudinal space-charge field E_z of the beam by what is called a “g-factor” model [6]. The

transverse variation of E_z is ignored, and the dependence on the longitudinal coordinate z is given in SI units by

$$E_z = -(g/4\pi\epsilon_0) d\bar{\rho}_b/dz, \quad (1)$$

where $\bar{\rho}_b$ is the beam line-charge density, ϵ_0 is the free-space permittivity, and g for a uniform-density beam with an average radius a centered in a beam pipe of radius R is

$$g = \ln(R^2/a^2). \quad (2)$$

In general, a will increase as the beam compresses, leading to a decrease in g . However, we assume that transverse focusing force increases appropriately with z , so that a and hence g remains fixed. For uniform density, the further assumption of a parabolic current variation in z leads to a space-charge field that varies linearly along the beam length. In the absence of longitudinal pressure, this linear field allows self-similar lengthwise compression of the profile. From the longitudinal envelope equation that results from these assumption, we find that the compression ratio C is given by

$$C = 1 + (8\bar{\rho}_0 g)^{-1} (\Delta v/v)^2, \quad (3)$$

where $\bar{\rho}_0$ is the initial value of the generalized beam perveance, written for a non-relativistic beam with ion charge state q and mass M as

$$\bar{\rho}_0 = (4\pi\epsilon_0)^{-1} 2qe\bar{\rho}_b/(Mv_z^3), \quad (4)$$

and $\Delta v/v$, referred to here as the “velocity tilt,” is shorthand for the initial difference between the longitudinal fluid velocity of the beam tail and that of the beam head, divided by the initial average longitudinal velocity v_z . The corresponding stagnation distance D is found to

$$D/L_{b0} = (\Delta v/v)^{-1} \{ (1-C^{-1}) + [(C-1)^{1/2} / (2C^{3/2})] \ln[2C + 2C^{1/2} (C-1)^{1/2} - 1] \}, \quad (5)$$

where we have scaled D to the initial beam length L_{b0} . By itself, the first term in Eq. (5) gives the propagation distance to reach minimum beam length for ballistic compression, while the second term corrects for space charge. For heavy-ion beams of interest for inertial fusion, perveance values range typically from 10^{-5} to 10^{-2} , and practical considerations are likely to limit the velocity tilt to less than 0.25 for driver-scale beams.

A second estimate of beam compression is obtained from numerical simulations using the code CIRCE [4], which divides a beam lengthwise into thin “slices,” each containing a constant amount of charge. The transverse dynamics of each slice is calculated from envelope equations for the radii and the centroid position in the two transverse directions, and the longitudinal motion of the each slice boundary is modeled by a Lagrangian fluid equation. The longitudinal space-charge field is

approximated here by a one of two methods: a generalized g -factor expression that accounts for variations of the beam density and the slice radii with z , and a Bessel series representation [6] that calculates the average E_z across slice boundaries from an exact analytic expression. The CIRCE dynamics model ignores longitudinal pressure, and it assumes that the transverse emittance of each slice is constant. Nonetheless, when benchmarked against the particle simulations, CIRCE is found to give reliable results, provided that the normalized emittance does not increase significantly and there is little longitudinal mixing. Because of its speed, CIRCE is a useful scoping code for lattice design.

These two fluid-like models are compared here with results from the 3-D electrostatic particle-in-cell (PIC) code WARP3d [5]. Although this code can model lattice elements from first principles, we specify hard-edge focusing quadrupoles and fringe-free acceleration fields to keep the lattice representation as close to the CIRCE model as possible, and we initialize WARP simulations with matched, uniform-density, longitudinally cold beams, like those assumed in CIRCE. For this initial work, we choose that the beam, the lattice, and the fields to be error free.

III. RESULTS

For all the simulations presented here, we use parameters appropriate for the Integrated Beam Experiment (IBX), a scaled experiment being considered at LBNL that would test all the major accelerator components for HIF. Although the peak current for this accelerator would be less than 1 A, the tune depression would match that in a HIF driver, and with proper design, nearly all aspects of drift-compression would be realistically modeled. The only physics that would not be correctly represented in IBX are image forces, stray-electron effects, and collisions between beam ions and with residual gas.

III.A Comparison of compression models

To allow a fair comparison of the three models, we choose a beam with a parabolic current profile and uniform density and emittance. The ion species used here is singly charged potassium (39 amu) with an energy of 17 MeV, so the 6.9-A peak current gives a maximum perveance of about 8.8×10^{-4} . We limit the scale of the simulation by choosing a 1.75-m initial beam length, corresponding to a 250-ns duration. The transverse emittance, which is relatively unimportant in these simulations, is 4.3 mm-mrad. A periodic FODO lattice is chosen to give an undepressed phase advance of 72° per 0.6-m lattice period and a initial average radius of 0.008 m.

The analytic compression model summarized by the higher fields found at the ends, accounting for the so




Figure 1. Plots of (a) stagnation distance and (b) compression ratio for uniform-density parabolic IBX-like beams with a range of initial velocity tilts. The overlaid “fixed g -factor” curve are calculated from the analytic expressions and by CIRCE. The “variable g -factor” and “Bessel series” are calculated using CIRCE with more-realistic models of the beam space-charge field. The points labeled “WARP” were calculated using WARP3d.

Eqs. (1) and (2) predicts a quadratic increase in the compression ratio with increasing initial velocity tilt $\Delta v/v$. Also, the compression length is predicted to increase from zero at for $\Delta v/v = 0$ and drop off after a single maximum roughly like $(\Delta v/v)^{-1}$. The curves labeled “fixed g -factor” in Fig. 1 show these analytic relations for the chosen 8.8×10^{-4} perveance, and qualitatively similar curves are found for other values.

Although the CIRCE simulations corroborate these qualitative predictions, significant quantitative differences are seen in the plots of Fig. 1 for different space-charge models. When the g -factor is artificially fixed in CIRCE, the code accurately reproduces the analytic curves, validating the analysis. However, the CIRCE curves calculated using the more accurate space-charge models show significantly higher compression ratios. Since the beam radius a is expected to increase as the beam compresses in a period focusing lattice, the g -factor in Eq. (2) and hence the space-charge field should decrease. Since it is this field that slows compression, both the higher compression ratios seen for larger tilts are plausible, as are the longer stagnation distances at larger $\Delta v/v$ values. For smaller initial tilts, the change in beam length and hence a becomes insignificant, and the stagnation distance is determined by how quickly the space-charge field at each end reverses the initial tilt. The fixed g -factor model underestimates the longitudinal space-charge field near the beam ends [6],

so the model predicts longer stagnation lengths than CIRCE for $\Delta v/v < 0.05$.

The corresponding WARP cases, plotted as points in Fig. 1, largely corroborate CIRCE model, with agreement being best for the Bessel-series E_z representation. This reduced compression appears to results from two features of the WARP3d model missing from the fluid models. First, both the analysis and the CIRCE model average the longitudinal space-charge field radially, whereas the Bessel-series representation in Ref. [6] shows that the field is peaked on axis. Since it is this on-axis value that affects the velocity of the beam ends, a larger value is expect to resist compression more successfully and stagnate the beam sooner. A second difference is the absence of longitudinal pressure in the fluid models considered here. WARP3d plots of the longitudinal phase space show some broadening of the distribution in v_z as a beam compresses, and this higher longitudinal temperature resists compression. More careful analysis of the WARP3d results is needed to determine the source of the longitudinal heating and whether either of these effects is large enough to account for the observed reduction in compression. If the lower compression seen in the WARP3d runs is verified, we would conclude that cold-fluid models cannot be relied upon for making quantitative predictions of unneutralized beam compression for HIF.

III.B Comparison of compression strategies

A somewhat more realistic beam and lattice are used to compare methods for shaping the beam velocity and current profiles before drift-compression. The beam is again 17-Mev potassium, but the 4-m beam length (1500-ns duration), 0.368-A peak current, and 10.3-mm-mrad emittance match existing equipment that could be used as an IBX front end. The perveance for this beam

the gaps. Longitudinal control fields, called “ears” are then added to the acceleration fields by running CIRCE with the longitudinal space-charge field artificially turned off. The space-charge field is calculated for the CIRCE beam as it passes through each gap and used to calculate the required ears. For the case here, we choose that the beam duration decreases parabolically with time and the average energy increases sufficiently that all pulses are non-negative. The resulting shaping fields,

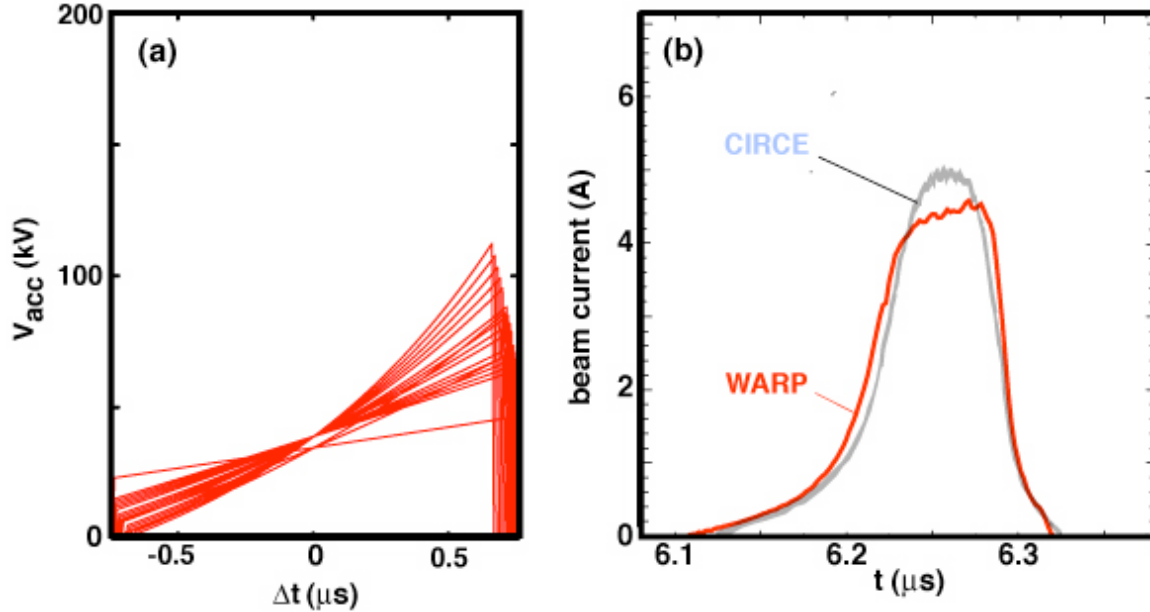


Figure 2. Plots of (a) beam-frame shaping fields and (b) the beam current profile at the stagnation point $z = 18.4$ m when the velocity tilt is imposed by approximately triangular pulses before drift-compression. The current curve labeled “CIRCE” is the same case calculated using the fluid-envelope code CIRCE.

is 6.7×10^{-4} , and the current profile has a 2-m flat top and drops to zero parabolically at the ends. A more complicated lattice is used for these cases. To prevent the outer edge of the beam from scraping on the beam-pipe wall, the field strength of focusing quadrupole magnets is gradually increased by about 35% over the second half of the 24-m lattice, and the beam-pipe radius increases stepwise from 0.03 m to 0.05 m over five lattice period (2.3 m) beginning at 18 m. Also, 0.05-m induction gaps are placed midway between quadrupoles in the first seventeen half-lattice periods to impose the beam-shaping electric fields.

The simplest method used here to impose a velocity tilt uses a previously published algorithm to control the beam energy and duration as functions of either time or the longitudinal position [7]. By assuming instantaneous acceleration in induction gaps and ignoring space-charge forces between gaps, it is straightforward to construct the trajectories of beams slices, and the resulting acceleration fields are then calculated as functions of time from the bends in slice trajectories at

shown in Fig. 2a shown as a function of time relative to the beam midpoint, are approximately triangular, and the corresponding ear fields, which retard ions at the beam head and slow those at the tail, reach a magnitude of about 4 kV at the beam ends. While this compression schedule achieves more than a ten-fold increase in the maximum current, as seen in Fig. 2b, the compression is not self-similar. Since no effort has been made to compensate for space charge as the beam compresses, nearly half the beam charge at stagnation lies outside the pulse length at half-maximum, so the compression ratio of the WARP3d beam is about 5:1. As in the idealized runs of Section III.A, the corresponding CIRCE run shows roughly the compression, although the model overestimates the peak current by about 10%.

To control the beam-end expansion seen in Fig. 1, we use the algorithm in Ref. [7] to set up a shaping schedule that first imposes a velocity tilt in eight gaps placed a full lattice period apart, then corrects and finally removes the tilt in thirteen more induction gaps spread over the remaining 19 m. Ears are added in all

gaps, augmenting the shaping field as tilt is applied and reducing it as the tilt is removed. Simulations indicate that this approach largely eliminates blow-off of the beam ends, so that less than 12% of the beam charge in the WARP3d profile is outside the full length at half-maximum. However, the shaping method is unsatisfactory for two important reasons. Due to the wider spacing of gaps, the maximum field strength is nearly double that of Fig. 2a, and the very short rise and fall times of the beam near stagnation require voltages for the ear pulses approaching 400 kV and roughly a 150-MHz response frequency. These requirements substantially exceed current pulser technology. Furthermore, we expect that the high voltage and frequency content of the ear fields would make the beams extremely sensitive to any timing errors,

of the shaping section. To calculate the shaping fields, we again ignore the longitudinal space-charge field and assume instantaneous acceleration in the gaps, so that beam slices have constant velocity between gaps. Slice trajectories are then piecewise linear lines in z - t space, with changes in slope at the induction gaps. We construct these trajectories to match the times and velocities of the input beam with the values required after shaping, and the required voltage in each gap is found from the change there in the trajectory slope. There are, of course, infinitely many possible trajectories. We have not yet developed a workable algorithm for optimizing the choice of trajectories, so here we choose the points where the slice trajectories intersect the gap by fitting a fifth-order polynomial to the initial and final data. Of course, a fourth-order

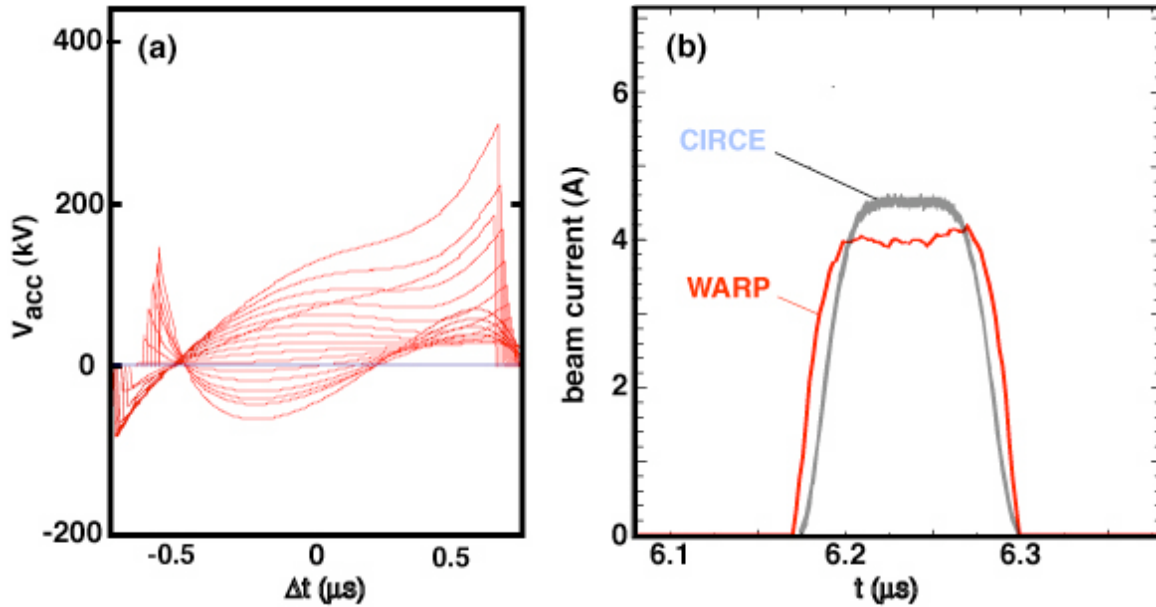


Figure 3. Plots of (a) beam-frame shaping fields and (b) the beam current profile at the stagnation point $z = 18.4$ m when the velocity and current profiles are tailored in a shaping lattice with 17 induction cells. This particular set of shaping fields was calculated using an unoptimized algorithm that allows energy to be removed and then added back to parts of the beam. The current curve labeled “CIRCE” is the same case calculated using the fluid-envelope code CIRCE.

although this sensitivity has not been tested. The other problem seen in WARP3d and CIRCE simulations is that the large ear fields launch a series of space charge waves that make compression far from self-similar.

A third approach to achieving controlled drift compression uses the beam-shaping algorithm similar to that of Ref. [2]. Using CIRCE, a flat-topped pulse with a profile similar to the input pulse but with a ten times the maximum current and a tenth the duration is first run backwards through the compression lattice, giving the velocity and current profiles that are needed at the end

polynomial can match the boundary conditions, but we use the additional degree of freedom to minimize the integral of the curvature along the trajectory. As before, ears are calculated in each gap from a CIRCE run, and the combined shaping and ear fields are tested in CIRCE by running the input beam through the entire lattice. The run is finally repeated using WARP3d to examine questions of phase-space dynamics, such as emittance growth.

We see from the calculated voltages in Fig. 3a that the choice of voltage waveforms is not optimal, since a

substantial amount of energy is first removed from slices near the head and then added back later. Also, the maximum voltage exceeds is nearly Fig. 3b shows that the WARP3d current profile at stagnation is a poor match to that used in the CIRCE run backward through the compression lattice. The compression ratio is about 10:1, and the compression is not self-similar, since irregular current spikes appear in the beam body and low current tails develop at both ends. We also find that the WARP3d compression length is 25% shorter than the CIRCE estimate. These differences suggest that CIRCE should not be used to calculate the profile needed after beam shaping. We plan to repeat this procedure using WARP3d rather than CIRCE for the backward run, and if emittance growth proves to be significant during drift, we can artificially reduce the beam emittance during the backward run. [virtually no particle loss, less than factor of two increase in transverse emittance, . The

V. Conclusions

We have begun using the 3-D electrostatic PIC code WARP3d to model unneutralized drift-compression of intense ion beams. Comparing WARP3d results for a linear head-to-tail velocity "tilt" with those of an analytic cold-fluid model and a more-realistic fluid-envelope simulation, we find that the particle model closely compression ratio and stagnation distance than the fluid models for a range of initial head-to-tail velocity tilts. In these cases, the transverse emittance in the particle simulation grows less than 10%, and the longitudinal thermal energy remains small compared with the electrostatic potential energy, so two of the assumptions in the fluid-envelope model are approximately valid. The analytic fluid model gives qualitatively correct scalings with beam perveance and initial tilt, but the predicted compression ratios are as much as 40% low for higher initial velocity tilts due to the very simple model of the longitudinal space-charge field that must be used to keep the analysis tractable.

Three methods have been tested for imposing a head-to-tail velocity tilt by time-dependent beam-shaping fields. Two methods prove to be unsatisfactory. Applying nearly triangular pulses allows the beam ends to expand excessively during compression. Using shaping fields that apply and then remove tilt before stagnation gives better control of the beam ends but places unreasonable demands on the voltage magnitude and time response of the pulsed power. The most successful of the methods constructs shaping fields that transform the input the velocity and current profiles into those found by allowing a compressed beam to expand as it moved backward through the drift lattice. A compression ratio of 10:1 is achieved, with adequate control of the final velocity and current profile, and the

transverse emittance grows by less than a factor of two. This method, while promising, poses a challenging optimization problem. Future work will focus on developing a more effective and robust algorithm to prepare beams for drift-compression.

Acknowledgments

This work was performed under the auspices of the US Department of Energy by University of California Lawrence Livermore National Laboratory and Lawrence Berkeley National Laboratory under Contracts No. W-7405-ENG-48 and DE-AC-3-76SF00098.

References

- [1] D. D.-M. Ho, S. T. Brandon, and E. P. Lee, *Part. Accel.* **35**, 15-42 (1991).
- [2] W. M. Sharp, A. Friedman, and D. P. Grote, "Manipulation of High-Current Pulses for Heavy-Ion Fusion," *AIP Conference Proceedings 391* (AIP, Woodbury, NY, 1997), pp. 27-35.
- [3] H. Qin, R. C. Davidson, J. J. Barnard, and E. P. Lee, "Drift Compression and Final Focus Options for Heavy Ion Fusion," *Proc. 2004 Heavy Ion Fusion Symposium*, 7-11 June 2004, Princeton, NJ (to be published).
- [4] W. M. Sharp, J. J. Barnard, D. P. Grote, S. M. Lund, and S. S. Yu, "Envelope Model of Beam Transport in ILSE" *AIP Conference Proceedings 297* (AIP, Woodbury, NY, 1994), pp. 540-548.
- [5] A. Friedman, D. P. Grote, and I. Haber, "Three Dimensional Particle Simulation of Heavy Ion Beams," *Phys. Fluids B* **4**, 2203-2210 (1992).
- [6] W. M. Sharp, D. A. Callahan, and D. P. Grote, "Effects of Longitudinal Space Charge in Beams for Heavy-Ion Fusion," *Fusion Eng. Design* **32-33**, 201 (1996).
- [7] W. M. Sharp and D. P. Grote, "Acceleration Schedules for a Recirculating Heavy-Ion Accelerator," *Phys. Rev. ST Accel. Beams* **5**, 094202 (2002).